Learning on the Web: A Content Literacy Perspective

John E. McEneaney

Department of Reading and Language Arts

Oakland University

Rochester, Michigan

United States

Published January 2000 in *Reading Online*, an electronic journal of the International Reading Association. Available at <http://www.readingonline.org/articles/mceneaney/>.

A Note about This Format

This paper is a reformatted version of an article originally created for presentation as a path-based web hypertext. The translation from hypertext to traditional linear format was not always straightforward. Web text that is impossible to interpret in a print version (e.g., passages about the use of navigational devices that relied on those devices being available) is not included here. In addition, since the text passages were developed as a set of discrete hypertext pages rather than as continuous text, transitions between paragraphs may be hard to follow. This version of the article was developed solely to highlight differences between traditional print and hypertext formats. Readers are

cautioned that it fails to capture dynamic qualities of the original presentation that can only be understood in an online reading environment.

Abstract

One of the fundamental principles of content literacy instruction is that supporting learning from text should involve a focus both on the content to be learned and on the processes students apply as they work to acquire, organize, and integrate that content (Readence, Bean, & Baldwin, 1995, pp. 9-10). The significance of this fundamental principle in terms of instructional outcomes is clearly related to the accessibility and complexity of the material to be learned, since more complex or less accessible text will require more sophisticated learning processes on the part of students. Unfortunately, the World Wide Web, although routinely touted as a breakthrough technology for educators, poses special difficulties with respect to the process focus advocated by content literacy experts (see, e.g., Anderson & Joerg, 1996; Edwards & Hardman, 1989; Nielsen, 1989). Fortunately, new dynamic capabilities of on-line reading environments and the tools they support should help software designers and educators to develop learning materials that can help readers avoid some of the problems with web-based content that have been documented in the literature.

The objective of this article is to briefly describe and, in the path-based version, to illustrate by example how new web technologies can be applied to assist readers both in

integrating content and in maintaining a process focus as they navigate complex expository text.

The Concept of a Path

One of the central ideas addressed in this article is that of a hypertext "path," a term that refers to a specific navigational sequence within a hypertext document. The concept of a path has significance for research in hypertext reading from two rather different perspectives. The *descriptive* perspective uses path data as a basis for empirical investigations of user movement in hypertext (Cardle, 1994; McEneaney, 1999; Smith, 1996). But the concept of path also has a *prescriptive* interpretation that is more immediately relevant to thinking about hypertext as a learning tool. Cognitive flexibility theory (Jacobson & Spiro, 1995; Spiro & Jehng, 1990), for instance, suggests that learning about a complex, ill-structured domain requires numerous carefully designed traversals (i.e., paths) across the terrain that defines that domain, and that different traversals yield different insights and understandings. Flexibility is thought to arise from the appreciation learners acquire for variability within the domain and their capacity to use this understanding to reconceptualize knowledge.

The important idea behind a path is that, in the absence of design, the navigation of an inexperienced reader may more closely resemble a random walk than a traversal likely to promote insight. Simply dropping students into complex hypertext may do little more than confuse and frustrate them. To the extent that we impose paths,

however, we also limit the ways our students can learn from and explore the materials we want them to read. The critical issue is whether the support we impose helps or hinders learning.

The Problem of Flexibility

If instructional materials are to promote flexibility, learners must first have knowledge that is effectively grounded -- they must understand, for example, that not all principles are created equal, and that concepts and principles are often related in subtle ways. In the absence of thoughtful instructional design, efforts to promote flexibility may only lead to disorganized thinking among students. This is one of the problems often cited concerning the web as a learning environment -- that it has so little intrinsic structure that the freedom it provides actually undermines learning in favor of a shallow browsing of material. A science teacher, for example, might want students to draw from many different science-oriented websites that have been developed by many different authors for many different audiences. This fragmentation of focus and voice will almost certainly have consequences for the students, whose task it is to integrate the information they gather. This means that instructors, either directly themselves or through supplemental materials they identify, will probably need to provide the instructional "glue" necessary to keep students focused on the task at hand so that the great diversity of the web truly becomes an asset rather than a handicap.

In response to problems related to hypertext navigation and users who become "lost in hyperspace," researchers and web-content developers have created a variety of powerful tools, many of them based on visualization techniques. Site maps are now commonly provided and there is evidence that users find them helpful for navigating and establishing a clearer idea of particular sites' organizational structure (Chen & Rada, 1996; Utting & Yankelovich, 1989). In larger networks where complete site maps are impractical, "fish-eye views" (Bartram, Ho, Dill, & Henigman, 1995; Furnas, 1986; Sarker & Brown, 1994), clustering techniques that organize nodes into meaningful groups (Gloor, 1991; Mukherjea, Foley, & Hudson, 1995), and a variety of other filtering and mapping techniques have been developed to assist in both creation and use of large-scale hypertext networks (e.g., Husemann, Petersen, Kanty, Kochs, & Hase, 1997; Kahn, 1999; Neves, 1997). On an even broader scale, web directories (increasingly referred to as "portals") have become important navigational landmarks on the web specifically because they provide structure.

But all these tools provide assistance only in the same sense that any map does, and users still have to figure out a great deal on their own. My purpose in this article is to argue, both by exposition and by example, that maps are *not* enough on the web. If we are serious about using this environment to support learning, we need to provide learners with substantially more guidance and help than is afforded by even the most ambitious site maps and web directories. We need to begin thinking about the web as a content literacy resource that requires the same attention to process as we recommend for traditional print materials.

A Continuum of Alternative Readings

As indicated on the abstract page of the online version of this article (see http://www.readingonline.org/articles/mceneaney/), this text can be read in several ways. One way is as a "traditional" web document that makes the reader responsible for blazing a trail across the content. In that version, readers select links on their own, thereby assuming responsibility for organizing the article's macrostructure. The obvious benefit of this approach is that whatever learning occurs has the potential to be very personal. The down side, however, is that empirical studies suggest that learners sometimes find it difficult to stay on track when navigational decision making is added to other demands. An alternative approach is to follow (closely or loosely) a predetermined path especially designed to preserve a particular perspective on the ideas presented (which may not be as clear when a reader uses his or her own path). A second version of this article presents the text (and a sample lesson) in path-based format.

The notion of alternative readings isn't exhausted, however, with just these two possibilities. It may, in fact, be more useful to think about the alternatives in terms of a *continuum* that ranges between two different poles that maximize either reader or author control.

Some may find reading either (or both) of the on-line versions of this article difficult. This may be because most of us are still novices when it comes to reading and writing in this format. My experience in writing hypertext to date leads me to conclude

that "hypercomposition" requires skills that differ substantially from those we traditionally associate with effective writing. Another aspect of writing hypertext I find interesting is that the boundaries between "writing" and "programming" seem to be very fluid. Some of the meaning I try to communicate in this article is grounded in the words you read. But some is also grounded in the scripts that power buttons and other navigational devices in the path-based version, determining *how* you read that document. The hidden text of the scripts can therefore have a very powerful influence on the way you read and understand the article.

Although it would be counterproductive to shackle readers permanently to a single perspective on the content presented, there are compelling pedagogical reasons to guide learners to greater or lesser degrees, depending on their needs and the instructional objectives. The central issue is one of designing appropriate levels of control so that learners can benefit from the experience and knowledge of others (especially teachers) while still finding room for their own unique insights and understanding. One of the most powerful aspects of the new online technologies is that these levels of control can be adjusted "on the fly." Learners who find the support offered too intrusive may be able to scale it back, and those who are in need of greater degrees of support can scale it up.

In the path-based version of this article, I attempted to incorporate degrees of control by providing both a path mechanism and links that allow readers to "step off" the path. Most of these links are internal to the document and provide a limited freedom to wander. Some, however, are external to the document (i.e. they are out on the web) and, as such, provide a substantial degree of freedom.

It is also relevant to note that the links that appear in the link panel at the bottom of the path-based version are dynamic -- they are created by a script at the time the page is loaded. The important point about links created by scripts is that they can be made contingent on a reader's prior choices. A common example of this sort of linking on the web occurs when one uses search engines (e.g. Yahoo or Altavista). If a search engine is used to find pages on a certain topic (say, travel), the advertisements that appear on the page of search results will almost invariably be related to the search terms, because they were accessed dynamically in response to those terms.

One way authors can use script-driven linking is to assure that readers visit pages that are deemed especially important. If a reader (for whatever reason) never loads a page that is deemed important, a script could gradually withdraw other links until the only remaining possibility is to load the important page that has not yet been visited. Moreover, this narrowing of alternatives could be either very explicit (links begin to disappear) or far more subtle (multiple links that appear to be independent actually direct the reader to the same page).

The Problem of Process

Although the web has certainly benefited from its novelty, there are some decided disadvantages to being the newest kid on the literacy block. We don't often spend time considering it, but as readers of print we rely on a wide range of "invisible" skills. We have learned how to hold a book or magazine and, as readers of English, to scan from left

to right; we have learned to use fonts, section headings, and page layout to distinguish important information. We also make use of conventions related to the structure of text itself, and our efforts as teachers of reading often revolve very directly around these conventions.

One significant problem with the web is that it is still so new that conventions have yet to emerge. This means readers need to devote energy and attention to processes that are usually automatic in traditional print. How do I find the index page in a complex document I've dropped into from a search engine? How do I know if a link takes me to a different section of the same document or to completely different material? How much confidence can I have that the material I am reading is factually accurate or authoritative in some way? How do I get back to where I started?

As users of traditional print, we face the same kinds of problems, but we have ready responses. Tables of contents are at the front. Distinct documents usually appear in separate chapters or books, articles or journals. We generally know something about the reputations associated with publishers of printed materials, and we have learned to rely on dog-ears, pencil marks in margins, and page numbers to help us track our progress. But we don't yet quite know all the ins and outs of "web literacy," and compounding the problem is the fact that the new technologies preserve much of the capabilities of the old ways while introducing new elements. The result is increased opportunities for complexity -- and as we increase complexity, we increase demands on those who seek to use the materials we produce.

The path-based version of this article provides an example. Although my intent in developing that reading environment was to explore a way to simplify the use of the web for delivering instructional content, the longer I worked on the path-based version, the clearer it became to me that I was, in fact, increasing the complexity of my manuscript as a result of my increased ambitions. In that version, I not only communicate content but I also braid commentary and analysis into the content threads I present. These broadened ambitions resulted in my trying to establish conventions that will probably be unfamiliar from most readers' prior experience with print -- and perhaps even from their prior web experience. Although these conventions might turn out to be useful, their unfamiliarity means that they introduce difficulties -- just as happens with any new learning or reading strategy when it is first encountered.

Since more complex materials will typically require more sophisticated learning processes on the part of students, we have a fairly compelling case for content area literacy instruction if we want to make use of web resources. But in the absence of convention it becomes less clear what we should be teaching. Until we know a bit more about the process itself, it isn't clear how we can support it.

Web-Based Instructional Content

Three related, but conceptually distinct, features of the web are particularly important in considering the web as instructional content: scope, accessibility, and structure. The scope of the web is enormous, with millions of distinct websites and hundreds of millions

of documents (Nielsen, 1999). How much freedom should we give our students in exploring this space? How many links are too many? With respect to instruction, scope introduces two specific problems. One is simply locating relevant material. That, of course, is the purpose of search engines and web directories, but as anyone who has used these tools knows, each engine and directory has its own particular strengths and weaknesses. The second problem has to do with the quality of materials on the web. Although "scope" can be understood to refer to the range of content areas represented, it can refer equally to the range of quality one finds within a single content area. Here, as before, the solution usually depends on a substantial investment of time and effort searching, reviewing, and selecting documents that support desired instructional outcomes.

The issue of web accessibility has at least two facets (Wresch, 1996). One has to do with acquiring the appropriate computer hardware and software and making the essential digital connection. In a classroom with a single web-capable computer it will not be easy to teach 20 or 30 students using web-based content. It may not be much easier, however, in a traditional school computer lab, where entire classes of students sit down for brief periods of work. Related to this is the issue of "over accessibility." The problem in this case is the ease with which students can follow links to other sites, regardless of whether they are related to the task at hand (Cruthirds & Hanna, 1996). This freedom of movement can make it difficult for students to stay on track and focused (Anderson & Joerg, 1996). Indeed, this is the instructional equivalent of the "lost in hyperspace"

problem that has been so widely studied and cited (see, e.g., Edwards & Hardman, 1989; Nielsen, 1989).

The second facet has to do with human "hardware" and "software" requirements related to the acquired skills and dispositions needed to ensure effective use. For example, if a university student in Ryazan, Russia, wants to write Javascript to power a website, she would have a significant advantage if she were also a reader of English. While it is not the case that being able to read and write English means one can code in Javascript, its mnemonic character (and that of most high-level programming languages) means one can make better guesses about what commands will do -- and, not surprisingly, the commands are easier to remember for English speakers.

Finally, as regards the web's structure -- it has very little. Any individual site is typically organized according to some system, but there is wide variation in the way different sites are organized. Taken as a whole, it would not be unfair to think of the web as an unsorted mountain of books rather than as a library. Search engines and directories manage to create pockets of structure, but in truth, apart from the structures imposed by domain names and individual authors, the data on the web are almost completely unsorted.

The lack of structure may well be the single most important problem from an instructional perspective. In a textbook, one expects both a clear "local" voice and a consistent organizational framework. While it might be reasonable to hold the same expectations for a single webpage or site, much of the power of the web arises from its capacity to provide access to many authors with many different perspectives. Under

these circumstances an expectation of a consistent organizational structure or voice is almost certain to be disappointed. If students are using web-based materials in their learning, there is a danger that unless careful consideration is given to instructional design, the forest will be lost among the trees since the conventions and conveniences of the traditional textbook no longer apply.

Some (particularly literary theorists such as George Landow, 1992, 1994) have suggested that this very characteristic liberates readers from traditional linear structures and thus brings text to a new level of potential. But as I have noted elsewhere, others have pointed out that users are sometimes paralyzed by possibilities and end up wandering about, trying to figure out where they are and where they want to be.

Although there are probably those who will disagree with my assessment, it seems to me that as educators we are obliged to address these problems if we wish to consider the web as a source of instructional content. While I would not argue with the idea that freedom to explore is a good thing and that learners should have an opportunity to create personal ways of knowing, it seems to me equally obvious that there is a time and place for imposed decision making, and that the web, probably more than most other reading environments, has something to gain from a judicious narrowing of possibilities.

Designing Web-Based Instructional Content

My general framework for thinking about the web as a learning tool centers on three concepts that go beyond those required for a general understanding of the web. These are *frame*, *path*, and *script*.

Frame on the web has a straightforward technical meaning. Most web browsers now support the display of multiple panels -- or "frames" -- within a browser window. Each frame in the display may operate independently or in close coordination with the others. Frames are now widely used for pages that display banners, navigation tools, and simultaneous access to multiple documents.

In the context of web-based instruction, what is most important about a frame is that it can provide a consistent local voice that will follow the student as she moves from site to site across the diversity of the web. Traditional text is usually authored by individuals who have a particular, identifiable voice and who create a linear structure that helps readers maintain a sense of continuity, but this is not generally true of the web. The virtue of a frame is that it can provide a kind of macrostructural "glue" that helps students integrate the diversity of the materials they find on the web by providing an organizing principle or voice.

Although frames can be distracting (see, e.g., Nielsen, 1996), I am convinced that they can be used to offer a valuable form of instructional support in on-line reading environments. They give web authors an opportunity to create the asides, sidebars, and glosses that have become increasingly common in printed textbooks. Moreover, the information provided in these asides can be tailored to the needs of individual readers. Web documents that begin by asking a reader about prior experience reading on line and

use the response as a basis for delivering experience-adjusted content can now be created relatively easily.

Paths are important both because they allow readers to take advantage of the experiences of other readers and because, whether they are shared or divergent, they can help establish a basis for more direct interaction and comparison. Hypertext authors and researchers have worked with the concept of paths for some time (see, e.g., Trigg, 1988; Zellweger, 1989), but recent developments in web authoring languages have dramatically simplified implementation of this feature. As has been stated, a path describes a specific traversal of nodes within a network of documents. In a sense, a path is like the "history" function on a browser that allows the user to move backward and forward node by node through previously visited sites.

What distinguishes a path from a browser history is that the former can be specified before the user begins a traversal and it can provide links to sites that have no connections built in by the original author(s). These are particularly important features from an instructional perspective. The capacity to define a path traversal allows an instructor to impose an organizational framework based both on the particular network of nodes selected and on a specific sequence in which they will be visited. The capacity to insert links ensures that an instructor will not be limited to those links that others have defined. Not only does this afford a greater degree of integration of content-related sites, it also provides a means to include new material as content on the web grows. (To encourage readers to explore the concept of a path a bit more directly, this article is presented in a path-based format at the *Reading Online* website.)

Note that although I tend to write of paths as if they required an author, an interesting variation is that of the "unauthored" path, an idea that may be clearer by analogy. Imagine a university that is planning a large, grassy quad. A landscape architect has been asked to lay out sidewalks to criss-cross that quad. The architect's response was to have the entire quad planted in grass and then to wait a year. At the end of the year, the architect returned to find that students and others who had navigated the quad had left their recommendation. The architect simply laid out sidewalks where the grass had been worn away. No one actually designed the layout of the walks -- it simply emerged from use.

The dynamic capabilities of web-based media mean that text can have the same kind of memory that the grass did. In effect, web text can remember who went where, how they got there, and how long they stayed. Moreover, this information can be used to dynamically adjust documents over time. When data suggest that users make frequent transitions from one page to another, that particular link can be highlighted while low frequency links can be minimized or dropped entirely -- the text can reconfigure itself as a result of use.

Finally, the concept of a script, or program that dynamically manages the presentation of text, is critically important because it provides authors a medium for directing the process of reading in a way that is not possible in traditional print. Scripts will be familiar to those who have developed materials in such software programs as Apple's HyperCard or Roger Wagner's HyperStudio. These programs allow teachers to create "cards" that link to one another in a way similar to sites and documents on the web.

Collections of cards are referred to as "stacks." But stacks can go beyond simple collections of text documents, simple images, or screen shots since they also support the use of scripts, short programs that can calculate some function or carry out a screen transformation when triggered by the user. The power of scripts lies in their capacity to support interactive presentation of material. Links can be determined dynamically based on the cards already visited or on some other feedback from the user.

With the development of Yahoo's Javascript and Microsoft's JScript and VB Script authoring languages, web pages can be made to include the same kind of interactivity that is possible in programs like HyperCard and HyperStudio. Moreover, these languages provide for better control over the text, images, and formatting that are so central to effective delivery of information through the web.

From an instructional perspective, one particularly important application of a script is to support a more interactive engagement of students with content. Scripts can be developed to guide students through pre-, during, and postreading activities, provide context-sensitive support, and even modify instructional content on the fly. Although frames provide the mechanism to display support and supplementary materials, it is the script running in the background that provides the plan. Other more subtle applications of scripts might involve creating user models like those employed in intelligent tutoring systems (Self, 1988), or using them as research tools to assist in collecting data about readers' navigation or comprehension.

One important application of a script, for example, involves its capacity to imbue text with "memory." Elsewhere I proposed the notion of "intentional" text -- text that has

an agenda of its own (McEneaney, 1997). A crucial element of intentional text is memory -- in tech-speak, this refers to persistent state variables that record "states" of the software while a reader is using it (e.g., software might record that a reader reviewed 12 pages in his reading session). The memory recorded by such scripts could be quite simple, limited to recording which pages were visited and for how long, or it could be quite rich if opportunities for interaction (and resulting data collection) had been provided to the reader during the session.

Conclusions

More than anything else, what distinguishes the web as an instructional resource from more traditional print-based materials is its capacity to support interactive presentation of content. Textbooks can provide appropriate glossing and metatext, and students are usually quite familiar with the "paths" laid out in printed material. In one sense, the web tools of frames and paths only serve to solve problems created by the new conventions of the on-line reading environment. The interactivity of the web, however, has a powerful compounding effect when coupled with frames and paths. As a result, these "fixes" assume new potential when they acquire dynamic capabilities.

I would argue that, from a content literacy perspective, the most significant benefit of new web technologies is to provide a basis for closer integration of content and process than has been possible in print media. Web-based materials can guide the learning process through direct manipulation of the reading environment based either on

predetermined schedules and plans or as a result of data collected while the learner is engaged with content. Moreover, even the content itself is subject to manipulation through the use of script-based linking and the means of delivery.

If we accept the premise that integration of process and content is desirable in promoting learning, and if we believe in the value of the web as an educational resource, I would argue that we must take a content literacy perspective if we are to make pedagogical sense out of this media amalgam. Although the use of web-based instructional resources poses special difficulties, it also provides powerful tools that can ameliorate or eliminate these problems and it introduces new possibilities that go far beyond what we have become accustomed to in print-based media.

References

Anderson, T.D., & Joerg, W.B. (1996). WWW to support classroom teaching. *Canadian Journal of Educational Communication*, *25*(1), 19-35.

Bartram, L., Ho, A., Dill, J., & Henigman, F. (1995). The continuous zoom: A constrained fisheye technique for viewing and navigating large information spaces. In *Proceedings, user interface and software technology '95* (pp. 207-215). New York: Association for Computing Machinery.

Cardle, N.T. (1994). *A HyperCard on Celtic history to assess navigability measures for hypertext* (Master of Science dissertation). Nottingham, UK: Institute of Information Technology, University of Nottingham.

Chen, C., & Rada, R. (1996). Interacting with hypertext: A meta-analysis of experimental studies. *Human-Computer Interaction, 11*, 125-156.

Cruthirds, J., & Hanna, M.S. (1996). *Programmed instruction and interactive multimedia: A third consideration* (rep. no. CS 509 364). Memphis, TN: Southern States Communication Association. (ERIC Document Reproduction Service No. ED 400 576)

Edwards, D.M., & Hardman, L. (1989). "Lost in hyperspace": Cognitive mapping and navigation in a hypertext environment. In R. McAleese (Ed.), *Hypertext: Theory into practice.* Oxford: Intellect Limited.

Furnas, G. (1986, April). Generalized fisheye views. In *Proceedings of CHI '86, Human Factors in Computing Systems* (pp. 16-23). New York: Association for Computing Machinery.

Gloor, P.A. (1991, December). Cybermap, yet another way of navigating in hyperspace. In *Proceedings of the Third Association for Computing Machinery Conference on Hypertext -- Hypertext '91* (pp. 107-121). New York: Association for Computing Machinery.

Husemann, H., Petersen, J., Kanty, C., Kochs, H., & Hase, P. (1997). A user adaptive navigation metaphor to connect and rate the coherence of terms and complex objects. In M. Bernstein, L. Carr, & K. Østerbye (Eds.), *Proceedings of the Eighth Association for Computing Machinery Conference on Hypertext -- Hypertext '97* (pp. 177-186). New York: Association for Computing Machinery.

Jacobson, M., & Spiro, R. (1995). Hypertext learning environments, cognitive flexibility, and the transfer of complex knowledge: An empirical investigation. *Journal of Educational Computing Research, 12*(4), 301-333.

Kahn, P. (1999, May). *Mapping web sites: Planning diagrams to site maps* [online seminar]. Providence, RI: Dynamic Diagrams. Available: <http://www.dynamicdiagrams.com/seminars/mapping/maptoc.htm>.

Landow, G. (1992). *Hypertext, the convergence of contemporary critical theory and technology*. Baltimore, MD: Johns Hopkins University Press. Chapter 1 available: <http://landow.stg.brown.edu/ht/contents.html>.

Landow, G. (1994). *Hyper/text/theory*. Baltimore, MD: Johns Hopkins University Press.

McEneaney, J.E. (1997, December 3). *Toward a post-critical theory of hypertext*. Paper presented at the 47th annual meeting of the National Reading Conference, Scottsdale, AZ. Available: <http://www.oakland.edu/~mceneane/nrc/conf97/nrc97ht5.pdf>.

McEneaney, J.E. (1999). Visualizing and assessing navigation in hypertext. In *Proceedings of the tenth ACM conference on hypertext: Returning to our diverse roots* (pp. 61-70). New York: Association for Computing Machinery. Available: <http://www.oakland.edu/~mceneane/nrc/conf99/symposium/paper2/metrics.pdf>

Mukherjea, S., Foley, J.D., & Hudson, S. (1995). Visualizing complex hypermedia networks through multiple hierarchical views. In *Conference proceedings on human factors in computing systems: CHI '95* (pp. 331-337). New York:

Association for Computing Machinery. Available: <http://www.acm.org/sigchi/chi95/Electronic/documnts/papers/sm_bdy.htm>.

Neves, F.D. (1997). The Aleph: A tool to spatially represent user knowledge about the WWW docuverse. In M. Bernstein, L. Carr, & K. Østerbye (Eds.), *Proceedings of the Eighth Association for Computing Machinery Conference on Hypertext -- Hypertext '97* (pp. 197-207). New York: Association for Computing Machinery.

Nielsen, J. (1989). The matters that really matter for hypertext usability. In *Proceedings of the Association for Computing Machinery Conference on Hypertext -- Hypertext '89* (pp. 239-248). New York: Association for Computing Machinery.

Nielsen, J. (1996, December). Why frames suck (most of the time). *Jakob Nielsen's Alertbox* [permanent content at useit.com: Jakob Nielsen's Website]. Available: <http://www.useit.com/alertbox/9612.html>.

Nielsen, J. (1999). User interface directions for the web. *Communications of the Association for Computing Machinery*, *42*(1), 65-72.

Readence, J.E., Bean, T.W., & Baldwin, R.S. (1995). *Content area literacy*. Dubuque, IA: Kendall/Hunt.

Sarker, M., & Brown, M.H. (1994, July 12). Graphical fisheye views. *Communications of the Association for Computing Machinery*, *37*, 73-84.

Self, J. (1988). *Artificial intelligence and human learning*. London: Chapman and Hall.

Smith, P.A. (1996). Towards a practical measure of hypertext usability. *Interacting with Computers*, *8*(4), 365-368.

Spiro, R., & Jehng, J. (1990). Cognitive flexibility and hypertext: Theory and technology for the nonlinear and multidimensional traversal of complex subject matter. In D. Nix & R. Spiro (Eds.), *Cognition, education, multimedia*. Hillsdale, NJ: Erlbaum.

Trigg, R. (1988). Guided tours and tabletops: Tools for communicating in a hypertext environment. *Association for Computing Machinery Transactions on Office Information Systems*, *6*(4), 398-414.

Utting, K., & Yankelovich, N. (1989). Context and orientation hypermedia networks. *Association for Computing Machinery Transactions on Office Information Systems*, *7*(1), 58-84.

Wresch, W. (1996). *Disconnected: Haves and have-nots in the information age*. New Brunswick, NJ: Rutgers University Press.

Zellweger, P.T. (1989). Scripted documents: A hypermedia path mechanism. In *Proceedings of the Association for Computing Machinery Conference on Hypertext -- Hypertext '89* (pp. 1-14). New York: Association for Computing Machinery.

Annotated Bibliography

General Reading

De Bra, P.M.E. (1996). *Hypermedia structures and systems*. Course 2L670, Eindhoven (The Netherlands) University of Technology. Available: <http://wwwis.win.tue.nl/2L670/static/index.html>.

Occasionally technical, but an excellent overview delivered in hypertext format.

Landow, G. (1992). *Hypertext: The convergence of contemporary critical theory and technology*. Baltimore, MD: Johns Hopkins University Press.

Outstanding introduction with a distinctively literary orientation.

Nielsen, J. (1993). *Multimedia and hypertext*. Boston: Academic.

A good nontechnical introduction with a computer science flavor.

Technical Reading

Annual Proceedings of the Association for Computing Machinery Conferences on Hypertext.

All are published by the Association for Computing Machinery (ACM) in New York. I have been able to obtain proceedings for '91 (San Antonio), '92 (Milan, Italy), '93 (Seattle), '94 (Edinburgh, Scotland), '96 (Washington, DC), and '97 (Southampton, UK). I highly recommend all of these proceedings for those with more technical interests. There have been others, but they are no longer in print.

Kommers, P.A.M., Grabinger, S., & Dunlap, J.C. (1996). *Hypermedia learning environments*. Mahwah, NJ: Erlbaum.

A semitechnical overview of hypertext from a design perspective.

Rouet, J., Levonen, J.J., Dillon, A., & Spiro, R.J. (1996). *Hypertext and cognition*. Mahwah, NJ: Erlbaum.

Empirical explorations of cognitive aspects of hypertext use.

Related Reading

Horton, W. (1994). *Designing and writing online documentation*. New York: Wiley.

A very user friendly introduction to technical communication in online formats.

Schneiderman, B. (1998). *Designing the user interface*. Reading, MA: Addison Wesley Longman.

A good starting point for a broader technical treatment. This volume also has a brief but useful chapter devoted to hypermedia and the web.